arXiv:1403.2401v2 [math.GT] 30 Jan 2015

GEOMETRIC GENERATORS FOR BRAID-LIKE GROUPS

DANIEL ALLCOCK AND TATHAGATA BASAK

Abstract. We study the problem of finding generators for the fundamental group G of a space of the following sort: one removes a family of complex hyperplanes from $\mathbb{C}^n$, or complex hyperbolic space $\mathbb{C}H^n$, or the Hermitian symmetric space for $\mathrm{O}(2,n)$, and then takes the quotient by a discrete group $P\Gamma$. The classical example is the braid group, but there are many similar "braid-like" groups that arise in topology and algebraic geometry. Our main result is that if $P\Gamma$ contains reflections in the hyperplanes nearest the basepoint, and these reflections satisfy a certain property, then G is generated by the analogues of the generators of the classical braid group. We apply this to obtain generators for G in a particular intricate example in $\mathbb{C}H^{13}$. The interest in this example comes from a conjectured relationship between this braid-like group and the monster simple group M, that gives geometric meaning to the generators and relations in the Conway-Simons presentation of $(M \times M) : 2$.

1. Introduction

The usual n-strand braid group of the plane was described by Fox and Neuwirth [FN] as the fundamental group of $\mathbb{C}^n$, minus the hyperplanes $x_i = x_j$, modulo the action of the group generated by the reflections across them (the symmetric group S_n). The idea is that a path $(x_1(t), \ldots, x_n(t))$ in this hyperplane complement corresponds to the braid whose strands are the graphs of the maps $t \mapsto x_i(t)$, regarded as curves in $[0,1] \times \mathbb{C}$. The removal of the hyperplanes corresponds to the condition that the strands do not meet each other. The fundamental group of the hyperplane complement is the pure braid group. Impure braids correspond to paths, not loops, in the hyperplane complement. But these paths become loops once we quotient by S_n, and one obtains the usual braid group.

The term "braid-like" in the title is meant to suggest groups that arise by this construction, generalizing the choices of $\mathbb{C}^n$ and this particular hyperplane arrangement. Artin groups [Br] and the braid groups of finite complex reflection groups [Be] are examples. The problem we address is: find generators for groups of this sort. We are mainly interested in the case that there are infinitely many hyperplanes, for example coming from hyperplane arrangements in complex hyperbolic space $\mathbb{C}H^n$. Our specific motivation is a conjecture relating the monster finite simple group to the braid-like group associated to a certain hyperplane arrangement in $\mathbb{C}H^{13}$. By our results and those of Heckman [H] and Heckman–Rieken [HR], this conjecture now seems approachable. We also suggest some applications to algebraic geometry.

Date: January 12, 2015.

2010 *Mathematics Subject Classification.* Primary: 57M05; Secondary: 20F36, 52C35, 32S22.

First author supported by NSF grant DMS-1101566.

The general setting is the following: Let X be complex Euclidean space, or complex hyperbolic space, or the Hermitian symmetric space for an orthogonal group $\mathrm{O}(2,n)$. Let $\mathcal{M}$ be a locally finite set of complex hyperplanes in X, $\mathcal{H}$ their union, and $P\Gamma \subseteq \operatorname{Isom} X$ a discrete group preserving $\mathcal{H}$. Let $a \in X - \mathcal{H}$. Then the associated "braid-like group" means the orbifold fundamental group

$$G_a := \pi_1^{\mathrm{orb}}\big((X-\mathcal{H})/P\Gamma, a\big)$$

See section 3 for the precise definition of this. In many cases, $P\Gamma$ acts freely on $X - \mathcal{H}$, so that the orbifold fundamental group is just the ordinary fundamental group.

Our first result describes generators for $\pi_1(X-\mathcal{H}, a)$. This is a subgroup of G_a, since $X-\mathcal{H}$ is an orbifold covering space of $(X-\mathcal{H})/P\Gamma$. For $H \in \mathcal{M}$ we define in section 2 a loop $\overline{aH}$ that travels from a to a point $c \in X-\mathcal{H}$ very near H, encircles H once, and then returns from c to a. We pronounce the notation "a loop H" or "a lasso H". See section 2 for details and a slight generalization (theorem 2.2) of the following result:

Theorem 1.1. *The loops $\overline{aH}$, with H varying over $\mathcal{M}$, generate $\pi_1(X-\mathcal{H}, a)$.* □

If a is generic enough then this follows easily from stratified Morse theory [GM]. But in our applications it is very important to take a non-generic, because choosing it to have large $P\Gamma$-stabilizer can greatly simplify the analysis of $\pi_1^{\mathrm{orb}}\big((X-\mathcal{H})/P\Gamma, a\big)$. So we prove theorem 1.1 with no genericity conditions on a. This lack of genericity complicates even the definition of $\overline{aH}$. For example, $\overline{aH}$ may encircle some hyperplanes other than H, and this difficulty cannot be avoided in any natural way (see remark 2.1). One might view theorem 1.1 as a first step toward a version of stratified Morse theory for non-generic basepoints.

Next we consider generators for the orbifold fundamental group G_a of $(X-\mathcal{H})/P\Gamma$. In our motivating examples, $P\Gamma$ is generated by complex reflections in the hyperplanes $H \in \mathcal{M}$. (A complex reflection means an isometry of finite order > 1 that pointwise fixes a hyperplane, called its mirror.) So suppose $H \in \mathcal{M}$ is the mirror of some complex reflection in $P\Gamma$. Then there is a "best" such reflection R_H, characterized by the following properties: every complex reflection in $P\Gamma$ with mirror H is a power of R_H, and R_H acts on the normal bundle of H by $\exp(2\pi i/n)$, where n is the order of R_H.

For each hyperplane $H \in \mathcal{M}$, we will define in section 3 an element $\mu_{a,H}$ of G_a which is the natural analogue of the standard generators for the classical braid group. Figure 1.1 illustrates the construction for the 3-strand braid group. (We have drawn the subset of $\mathbb{R}^3 \subseteq \mathbb{C}^3$ with coordinate sum 0, and our paths lie in this $\mathbb{R}^2$ except where they dodge the hyperplanes.) Recall that the definition of $\overline{aH}$ referred to a point $c \in X-\mathcal{H}$ very near H, and a circle around H based at c. We define $\mu_{a,H}$ to go from a to c as before, then along the portion of this circle from c to $R_H(c)$, then along the R_H-image of the inverse of the path from a to c. This is a path in $X-\mathcal{H}$, not a loop. But R_H sends its beginning point to its end point, so we have specified a loop in $(X-\mathcal{H})/P\Gamma$. So we may regard $\mu_{a,H}$ as an element of G_a. (Because a may have nontrivial stabilizer, properly speaking we must record the ordered pair $(\mu_{a,H}, R_H)$ rather than just $\mu_{a,H}$; see section 3 for background on the orbifold fundamental group.)

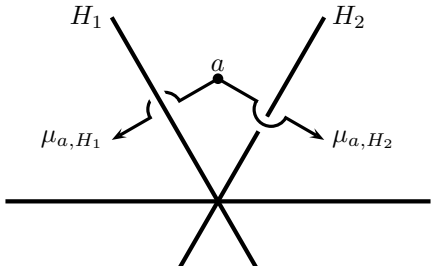

FIGURE 1.1. The standard generators for the 3-strand braid group, in terms of the A_2 hyperplane arrangement.

Referring again to figure 1.1, generation of the 3-strand braid group $\pi_1^{\mathrm{orb}}\big((\mathbb{C}^2 - \mathcal{H})/S_3, a\big)$ requires only the loops $\mu_{a,H}$ for the hyperplanes H closest to the basepoint a. This is proven in [FN] for the n-strand braid group. For this to hold, one should choose a as we did here, having the same distance to every facet of the Weyl chamber. Our next theorem shows that the same holds in our more general situation, provided that a is chosen so that there are "enough" mirrors closest to it. This analogy with the braid group is the source of the term "geometric generators" in our title.

Theorem 1.2. *Suppose $\mathcal{C} \subseteq \mathcal{M}$ are the hyperplanes closest to a, and that the complex reflections R_C generate $P\Gamma$, where C varies over $\mathcal{C}$. Suppose that for each $H \in \mathcal{M} - \mathcal{C}$, some power of some R_C moves a closer to p, where p is the point of H closest to a. Then the loops $(\mu_{a,C}, R_C)$ generate $G_a = \pi_1^{\mathrm{orb}}\big((X - \mathcal{H})/P\Gamma, a\big)$.*

The hypothesis of being able to move a closer to the various p's appears very hard to check in practice. But it can be done in some nontrivial cases; its verification (slightly weakened) for our motivating example occupies most of section 5.

This motivating example is the setting for a conjectural relation between a particular braid-like group and the monster simple group M. Here are minimal details; see section 4 for more background and [A4] for full details. We take X to be complex hyperbolic space $\mathbb{C}H^{13}$, and $P\Gamma$ to be a particular finite-covolume discrete subgroup of $\operatorname{Aut}\mathbb{C}H^{13} = PU(13,1)$, generated by complex reflections of order 3. We take $\mathcal{M}$ to be the set of mirrors of the complex reflections in $P\Gamma$. It turns out that any two mirrors are $P\Gamma$-equivalent, so the image of $\mathcal{H}$ in $X/P\Gamma$ is irreducible. The positively-oriented boundary, of a small disk in $X/P\Gamma$ transverse to a generic point of this image, determines a conjugacy class in $\pi_1^{\mathrm{orb}}\big((X - \mathcal{H})/P\Gamma\big)$. Following knot theory terminology, we call the elements of this conjugacy class *meridians*.

Conjecture 1.3 ([A4])**.** *The quotient of $\pi_1^{\mathrm{orb}}\big((X-\mathcal{H})/P\Gamma\big)$, by the normal subgroup generated by the squares of the meridians, is the semidirect product of $M \times M$ by $\mathbb{Z}/2$, where M is the monster simple group and $\mathbb{Z}/2$ exchanges the factors in the obvious way.*

Presumably, any proof of this will require generators and relations for $\pi_1^{\mathrm{orb}}\big((X - \mathcal{H})/P\Gamma\big)$, which is the motivation for the current paper. In [Ba1] the second author found a point $\tau \in X - \mathcal{H}$ (called $\bar{\rho}$ there), such that the set $\mathcal{C}$ of mirrors closest to τ has size 26, and showed that their complex reflections generate $P\Gamma$. Describing τ precisely requires some preparation, so we refer to section 4 for details and just

mention that τ has nontrivial $P\Gamma$-stabilizer. Therefore the corresponding meridians are ordered pairs $(\mu_{\tau,C}, R_C)$ rather than just bare paths $\mu_{\tau,C}$. Taking τ as our basepoint, we announce the following result, which we regard as a significant step toward conjecture 1.3.

Theorem 1.4. *The* 26 *meridians* $(\mu_{\tau,C}, R_C)$*, with* C *varying over the* 26 *mirrors of* $\mathcal{M}$ *closest to* τ*, generate* $\pi_1^{\text{orb}}\big((X-\mathcal{H})/P\Gamma, \tau\big)$. □

We wish this were a corollary of theorem 1.2. Unfortunately the hypothesis there about being able to move τ closer to the various $p \in H \in \mathcal{M}$ fails badly. Instead, the proof goes as follows. First, in section 5 we prove theorem 1.5 below, which is the analogue of theorem 1.4 with a different basepoint ρ in place of τ. For this basepoint, one can (almost) verify the hypothesis of theorem 1.2 about being able to move ρ closer to the various p's. Considerable calculation is required, together with extra work dealing with the fact that this hypothesis almost holds but not quite.

Then, given theorem 1.5, one joins τ and ρ by a path and uses it to identify the fundamental groups based at these points. One can then prove theorem 1.4 by studying how these groups' generators are related. The argument is delicate, of more specialized interest, and has little in common with the ideas in this paper. Therefore it will appear separately.

After stating theorem 1.5 we will explain our real reason for preferring τ over ρ as a basepoint. A minor additional reason is that ρ is not a point of $\mathbb{C}H^{13}$. Rather, it is a cusp of $P\Gamma$, and in particular lies in the sphere at infinity $\partial\mathbb{C}H^{13}$. This complicates things in two ways. First, there are infinitely many mirrors "closest" to ρ, indexed by the elements of a certain 25-dimensional integral Heisenberg group. And second, the definition of the meridians "based at ρ" requires more care. This leads to theorem 1.5 giving an infinite generating set, consisting of paths that are more complicated than those of theorem 1.4.

One can work through these complications as follows. As we did for τ, we refer to section 4 for the precise definition of ρ. All we need for now is that it is a cusp of $P\Gamma$ and that there is a closed horoball A centered at ρ that misses $\mathcal{H}$ (lemma 4.1). We choose any basepoint a inside A. We call the mirrors that come closest to A the "Leech mirrors". The name comes from the fact that they are indexed by the elements of (a central extension of) the complex Leech lattice Λ; in particular there are infinitely many of them. If C is a Leech mirror, let $b_C \in A$ be the point of A nearest it. Then $\mu_{a,A,C}$ is defined to be the geodesic $\overline{ab_C} \subseteq A$ followed by $\mu_{b_C,C}$ followed by $R_C(\overline{b_C a})$. See figure 3.1 for a picture. These are meridians in the sense of conjecture 1.3, and we call them the Leech meridians. As before, because a may have nontrivial $P\Gamma$-stabilizer, the meridian associated to C is really the ordered pair $(\mu_{a,A,C}, R_C)$ rather than just the bare path $\mu_{a,A,C}$.

Theorem 1.5. *The orbifold fundamental group* $\pi_1^{\text{orb}}\big((X-\mathcal{H})/P\Gamma, a\big)$ *is generated by the Leech meridians, that is, by the loops* $(\mu_{a,A,C}, R_C)$ *with* C *varying over the (infinitely many) mirrors closest to* ρ*.*

We promised to explain the real reason we prefer theorem 1.4 to theorem 1.5, that is, why we prefer the basepoint to be τ rather than ρ. It is because the 26 meridians of theorem 1.4 are closely related to the coincidences that motivated conjecture 1.3. In particular, by [Ba2] they satisfy the braid and commutation relations specified by the incidence graph of the 13 points and 13 lines of $P^2(\mathbb{F}_3)$.

That is, two generators x, y commute ($xy = yx$) or braid ($xyx = yxy$) according to whether the corresponding nodes of this graph are unjoined or joined. We call the abstract group with 26 generators, subject to these relations, the Artin group of $P^2(\mathbb{F}_3)$.

There are a family of presentations of the "bimonster" $(M \times M) \rtimes \mathbb{Z}/2$ as a quotient of this Artin group, due to Conway, Ivanov, Norton, Simons and Soicher in various combinations. All of them impose the relations that the generators have order 2, which yields the (infinite) Coxeter group whose Coxeter diagram is the same incidence graph. Modding out by the squares of meridians in conjecture 1.3 corresponds to taking this quotient. There are several different ways to specify additional relations that collapse this Coxeter group to the bimonster. The most natural one for our purposes seems to be the "$\widetilde{A}_{11}$ deflation" relations of [CS], because these have a good geometric interpretation in terms of the $\mu_{\tau,C}$'s and a certain copy of $\mathbb{C}H^9$ in $\mathbb{C}H^{13}$. See [H] and [HR] for more details.

We hope that our techniques will be useful more generally. For example, they might be used to give generators for the fundamental group of the moduli space of Enriques surfaces. Briefly, this is the quotient of the Hermitian symmetric space for $\mathrm{O}(2, 10)$, minus a hyperplane arrangement, by a certain discrete group. See [Na] for the original result and [A3] for a simpler description of the arrangement. The symmetric space has two orbits of 1-dimensional cusps, one of which misses all the hyperplanes. Taking this as the base "point", the hyperplanes nearest it are analogues of the Leech mirrors. It seems reasonable to hope that the meridians associated to these mirrors generate the orbifold fundamental group.

There are many spaces in algebraic geometry with a description $(X - \mathcal{H})/P\Gamma$ of the sort we have studied. For example, the discriminant complements of many hypersurface singularities [L1][L2], the moduli spaces of del Pezzo surfaces [ACT1] [Ko1][HL], the moduli space of curves of genus four [Ko2], the moduli spaces of smooth cubic threefolds [ACT2][LS] and fourfolds [L4], and the moduli spaces of lattice-polarized K3 surfaces [Ni][D]. The orbifold fundamental groups of these spaces are "braid-like" in the sense of this paper, and we hope that our methods will be useful in understanding them.

The paper is organized as follows. In section 2 we study the fundamental group $\pi_1(X - \mathcal{H}, a)$, in particular proving theorem 1.1. The proof relies on van Kampen's theorem. In section 3 we study $\pi_1^{\mathrm{orb}}\big((X - \mathcal{H})/P\Gamma, a\big)$, in particular proving theorem 1.2. The core of that proof is lemma 3.1, which is more general than needed for theorem 1.2. The extra generality is needed for our application to $\mathbb{C}H^{13}$. Section 4 gives background on complex hyperbolic space and the particular hyperplane arrangement referred to in conjecture 1.3 and theorems 1.4–1.5. Finally, section 5 proves theorem 1.5. Most of the proof consists of tricky calculations verifying the hypothesis of theorem 1.2 that the basepoint can be moved closer to the various points $p \in H$. In a few cases this is not possible, so we have to do additional work.

The first author is very grateful to RIMS at Kyoto University for its hospitality during two extended visits while working on this paper. The second author would like to thank Kavli-IPMU at University of Tokyo for their hospitality during two one-month visits while working on this paper.

2. Loops in arrangement complements

For the rest of the paper we fix $X =$ one of three spaces, $\mathcal{M} =$ a locally finite set of hyperplanes in X, and $\mathcal{H} =$ their union. The precise assumption on X is that it is complex affine space with its Euclidean metric, or complex hyperbolic space, or the Hermitian symmetric space for $\mathrm{O}(2,n)$. To understand the general machinery in this section and the next, it is enough to think about the affine case. In our application in section 5 we specialize to the case that X is complex hyperbolic 13-space; for background see section 4. Most of the other potential applications mentioned in the introduction would use the $\mathrm{O}(2,n)$ case.

We will freely use a few standard properties of X: it is contractible, and its natural Riemannian metric is complete and has nonpositive sectional curvature. By [BH, Theorems II.1A.6 and II.4.1], another way to state these properties is that X is a complete CAT(0) space. We will also use the usual notions of complex lines and complex hyperplanes in X, both of which are totally geodesic.

For $b, c \in X$ we write $\overline{bc}$ for the geodesic segment from b to c. Now suppose $b, c \notin \mathcal{H}$. It may happen that $\overline{bc}$ meets $\mathcal{H}$, so we will define a perturbation $\widehat{bc}$ of $\overline{bc}$ in the obvious way. The notation may be pronounced "b dodge c" or "b detour c". We write $\overline{bc}^{\mathbb{C}}$ for the complex line containing $\overline{bc}$. By the local finiteness of $\mathcal{M}$, $\overline{bc}^{\mathbb{C}} \cap \mathcal{H}$ is a discrete set. Consider the path got from $\overline{bc}$ by using positively oriented semicircular detours in $\overline{bc}^{\mathbb{C}}$, around the points of $\overline{bc} \cap \mathcal{H}$, in place of the corresponding segments of $\overline{bc}$. After taking the radius of these detours small enough, the construction makes sense and the resulting homotopy class in $X - \mathcal{H}$ (rel endpoints) is radius-independent. This homotopy class is what we mean by $\widehat{bc}$.

At times we will need to speak of the "restriction of $\mathcal{H}$ at p", where p is a point of X. So we write $\mathcal{M}_p$ for the set of hyperplanes in $\mathcal{M}$ that contain p, and $\mathcal{H}_p$ for their union.

Now suppose $b \in X - \mathcal{H}$ and $H \in \mathcal{M}$. We will define a homotopy class $\overline{bH}^{\circ}$ of loops in $X - \mathcal{H}$ based at b; the notation can be pronounced "b loop H" or "b lasso H". We write p for the point of H nearest b. It exists and is unique by H's convexity and X's nonpositive curvature [BH, Prop. II.2.4]. Let U be a ball around p that is small enough that $U \cap \mathcal{H} = U \cap \mathcal{H}_p$, and let c be a point of $(\overline{bp} \cap U) - \{p\}$. Consider the circular loop in $\overline{bp}^{\mathbb{C}}$ centered at p, based at c, and traveling once around p in the positive direction. It misses $\mathcal{H}$, because under the exponential map $T_pX \to X$, the elements of $\mathcal{M}_p$ correspond to complex hyperplanes in T_pX, while $\overline{bp}^{\mathbb{C}}$ corresponds to a complex line. And the line misses the hyperplanes except at 0, because $b \notin \mathcal{H}$. Finally, $\overline{bH}^{\circ}$ means $\widehat{bc}$ followed by this circular loop, followed by reverse($\widehat{bc}$).

Remark 2.1 (Caution in the non-generic case). This definition has some possibly unexpected behavior when b is not generic. For example, take $\mathcal{M}$ to be the A_2 arrangement in $X = \mathbb{C}^2$, let H be one of the three hyperplanes, and take $b \in X - \mathcal{H}$ orthogonal to H. See figure 1.1, with that figure's a being b, and the unlabeled hyperplane being H. It is easy to see that $\overline{bH}^{\circ}$ encircles *all three hyperplanes*, not just H. Furthermore, this phenomenon cannot be avoided by any procedure that respects symmetry. To explain this we note that $\pi_1(X - \mathcal{H}, b) \cong \mathbb{Z} \times F_2$, where the first factor is generated by $\overline{bH}^{\circ}$ and the second is free on $\overline{bH_1}^{\circ}$ and $\overline{bH_2}^{\circ}$, and H_1 and H_2 are the other two hyperplanes. Let f be the isometry of X that fixes b and negates H. It exchanges H_1 and H_2. So f's action on $\pi_1(X - \mathcal{H}, b)$ fixes $\overline{bH}^{\circ}$ and

swaps the other two generators. It follows that the group of fixed points of f in $\pi_1(X - \mathcal{H}, b)$ is just the first factor $\mathbb{Z}$. So any symmetry-respecting definition of $\overline{bH}^{\circ}$ must give some power of our definition.

The main result of this section, theorem 2.2, shows that the various $\overline{bH}^{\circ}$ generate $\pi_1(X - \mathcal{H}, b)$. But for our applications to $\mathbb{C}H^{13}$ in section 5, it will be useful to formulate the fundamental group with a "fat basepoint" A in place of b. This is because we will want to choose our basepoint to be a cusp of a finite-covolume discrete subgroup of $\mathrm{PU}(13,1) = \mathrm{Aut}\,\mathbb{C}H^{13}$. Strictly speaking this is not possible, since a cusp is not a point of $\mathbb{C}H^{13}$. So we will use a closed horoball A centered at that boundary point in place of a basepoint. For purposes of understanding the current section, the reader may take A to be a point.

Our assumptions so far are that X is one of three spaces, $\mathcal{M}$ is a locally finite hyperplane arrangement, and $\mathcal{H}$ is the union of the hyperplanes. Henceforth we also assume that A is a nonempty closed convex subset of X, disjoint from $\mathcal{H}$. To avoid some minor technical issues, we assume two more properties, both automatic when A is a point. First, for every $H \in \mathcal{M}$, there is a unique point of A closest to H. (This holds if A is strictly convex, by the argument used for [BH, Prop. II.2.4].) Second, some group of isometries of X, preserving $\mathcal{M}$ and A, acts cocompactly on the boundary ∂A. (This holds in our application to $\mathbb{C}H^{13}$ because the stabilizer of a cusp acts cocompactly on any horosphere centered there.)

We will think of the open r-neighborhood B_r of A, for some $r > 0$ as being "like" an open ball. If A is a point then of course B_r actually is one. In any case, the convexity of A implies that of B_r by [BH, Cor. II.2.5] and the remark preceding it.

Because $A - \mathcal{H} = A$ is simply connected (even contractible), the fundamental groups of $X - \mathcal{H}$ based at any two points of A are canonically identified. So we write just $\pi_1(X - \mathcal{H}, A)$ for $\pi_1(X - \mathcal{H}, a)$, where a is any point of A. If $c \in X$ then we define $\overline{Ac}$ as $\overline{bc}$, where b is the point of A nearest c. If $c \notin \mathcal{H}$ then we also define $\widehat{\overline{Ac}}$ as $\widehat{\overline{bc}}$. Similarly, if $H \in \mathcal{M}$ then we define $\overline{AH}^{\circ}$ as $\overline{bH}^{\circ}$, where b is the point of A closest to H. We sometimes write $\overline{b, c}$ and $\widehat{\overline{b, c}}$ and $\overline{b, H}^{\circ}$ for $\overline{bc}$ and $\widehat{\overline{bc}}$ and $\overline{bH}^{\circ}$, and similarly for $\overline{A, c}$ and $\widehat{\overline{A, c}}$ and $\overline{A, H}^{\circ}$.

Theorem 2.2 (π_1 of a ball-like set minus hyperplanes)**.** *Let B_r be the open r-neighborhood of A, where $r \in (0, \infty]$. Then $\pi_1(B_r - \mathcal{H}, A)$ is generated by the $\overline{AH}^{\circ}$'s for which $d(A, H) < r$.*

If A is compact, for example $A = \{a\}$, then this gives a finite number of generators for $\pi_1(B_r - \mathcal{H})$. But if A is non-compact then the number of generators may be infinite. This happens in theorem 1.5, where A is a horoball in $\mathbb{C}H^{13}$. The rest of the section is devoted to the proof of theorem 2.2, beginning with two lemmas.

Lemma 2.3 (π_1 of $\mathbb{C}^n$ minus hyperplanes through the origin)**.** *Suppose X is complex Euclidean space, every $H \in \mathcal{M}$ contains the origin 0, and $c \in X - \mathcal{H}$. Write $\frac{1}{2}X$ for the open halfspace of X that contains c and is bounded by the real orthogonal complement to $\overline{c0}$. (In the trivial case $\mathcal{M} = \emptyset$ we also assume $c \neq 0$, so that $\frac{1}{2}X$ is defined.)*

(1) *If c is not orthogonal to any element of $\mathcal{M}$, then $\pi_1(X - \mathcal{H}, c)$ is generated by $\pi_1(\frac{1}{2}X - \mathcal{H}, c)$.*

(2) *If c is orthogonal to some $H \in \mathcal{M}$, then $\pi_1(X - \mathcal{H}, c)$ is generated by $\pi_1(\frac{1}{2}X - \mathcal{H}, c)$ together with any element of $\pi_1(X - \mathcal{H}, c)$ having linking number ± 1 with H, for example $\overline{c\circ H}$.*

Proof. (2) Write H' for the translate of H containing c. Every point of $X - \mathcal{H}$ is a nonzero scalar multiple of a unique point of $H' - \mathcal{H}$. It follows that $X - \mathcal{H}$ is the topological product of $H' - \mathcal{H} \subseteq \frac{1}{2}X - \mathcal{H}$ with $\mathbb{C} - \{0\}$. The map $\pi_1(X - \mathcal{H}, c) \to \mathbb{Z}$ corresponding to the projection to the second factor is the linking number with H. All that remains to prove is that $\overline{c\circ H}$ has linking number 1 with H. In fact more is true: essentially by definition, this loop generates the fundamental group of the factor $\mathbb{C} - \{0\}$.

(1) We define H as the complex hyperplane through 0 that is orthogonal to $\overline{c0}$. We apply the previous paragraph to $\mathcal{M}' = \mathcal{M} \cup \{H\}$ and $\mathcal{H}' = \mathcal{H} \cup H$. Using $\frac{1}{2}X - \mathcal{H} = \frac{1}{2}X - \mathcal{H}'$ yields

$$(*) \qquad \pi_1(X - \mathcal{H}', c) = \pi_1(\tfrac{1}{2}X - \mathcal{H}', c) \times \langle \overline{c\circ H} \rangle = \pi_1(\tfrac{1}{2}X - \mathcal{H}, c) \times \langle \overline{c\circ H} \rangle.$$

Our goal is to show that the first factor on the right surjects to $\pi_1(X - \mathcal{H}, c)$. Let γ be any element of $\pi_1(X - \mathcal{H}', c)$ that is freely homotopic to the boundary of a small disk transverse to H at a generic point of H. It dies under the natural map $\pi_1(X - \mathcal{H}', c) \to \pi_1(X - \mathcal{H}, c)$. Because γ has linking number ± 1 with H, the product decomposition (*) shows that every element of $\pi_1(X - \mathcal{H}', c)$ can be written as a power of γ times an element of $\pi_1(\frac{1}{2}X - \mathcal{H}, c)$. It is standard that $\pi_1(X - \mathcal{H}', c) \to \pi_1(X - \mathcal{H}, c)$ is surjective. (Take any loop in $X - \mathcal{H}$, perturb it to miss H, and then regard it as a loop in $X - \mathcal{H}'$.) Since this map kills γ, it must send the subgroup $\pi_1(\frac{1}{2}X - \mathcal{H}, c)$ of $\pi_1(X - \mathcal{H}', c)$ surjectively to $\pi_1(X - \mathcal{H}, c)$. □

Lemma 2.4 (π_1 of a ball-like set with a bump, minus hyperplanes)**.** *Suppose $r > 0$, B is the open r-neighborhood of A, and $p \in \partial B$. Assume U is any open ball centered at p, small enough that $U \cap \mathcal{H} = U \cap \mathcal{H}_p$.*

(1) *If no $H \in \mathcal{M}_p$ is orthogonal to $\overline{Ap}$, then $\pi_1\big((B \cup U) - \mathcal{H}, A\big)$ is generated by the image of $\pi_1(B - \mathcal{H}, A)$.*
(2) *If some $H \in \mathcal{M}_p$ is orthogonal to $\overline{Ap}$, then $\pi_1\big((B \cup U) - \mathcal{H}, A\big)$ is generated by the image of $\pi_1(B - \mathcal{H}, A)$, together with any loop of the following form $\alpha\lambda\alpha^{-1}$, for example $\overline{A\circ H}$. Here α is a path in $B - \mathcal{H}$ from A to a point of $(B \cap U) - \mathcal{H}$ and λ is a loop in $U - \mathcal{H}$, based at that point and having linking number ± 1 with H.*

Proof. For uniformity, in case (1) we choose some path α in $B - \mathcal{H}$ beginning in A and ending in $(B \cap U) - \mathcal{H}$. In both cases we write c for the final endpoint of α; without loss of generality we may suppose $c \in \overline{Ap} - \{p\}$. Van Kampen's theorem shows that $\pi_1\big((B \cup U) - \mathcal{H}, c\big)$ is generated by the images of $\pi_1(B - \mathcal{H}, c)$ and $\pi_1(U - \mathcal{H}, c)$. We claim that $\pi_1(U - \mathcal{H}, c)$ is generated by the image of $\pi_1\big((B \cap U) - \mathcal{H}, c\big)$, supplemented in case (2) by λ.

Assuming this, we move the basepoint from c into A along reverse(α). This identifies $\pi_1(B - \mathcal{H}, c)$ with $\pi_1(B - \mathcal{H}, A)$, λ with $\alpha\lambda\alpha^{-1}$, and the elements of $\pi_1\big((B \cap U) - \mathcal{H}, c\big)$ with certain loops in $B - \mathcal{H}$ based in A. It follows that $\pi_1\big((B \cup U) - \mathcal{H}, A\big)$ is generated by the image of $\pi_1(B - \mathcal{H}, A)$, supplemented in case (2) by $\alpha\lambda\alpha^{-1}$. This is the statement of the theorem.

So it suffices to prove the claim. We transfer this to a problem in the tangent space T_pX by the exponential map and its inverse (written log). So we must show that $\pi_1(\log U - \log \mathcal{H}_p, \log c)$ is generated by the image of $\pi_1\big(\log(B \cap U) - \log \mathcal{H}_p, \log c\big)$, supplemented in case (2) by $\log \lambda$. The key to this is that the vertical arrows in the following commutative diagram are homotopy equivalences.

$$\begin{array}{ccc} \log(B \cap U) - \log \mathcal{H}_p & \longrightarrow & \log U - \log \mathcal{H}_p \\ \downarrow & & \downarrow \\ \frac{1}{2}T_pX - \log \mathcal{H}_p & \longrightarrow & T_pX - \log \mathcal{H}_p \end{array}$$

Here $\frac{1}{2}T_pX$ is as in lemma 2.3: the open halfspace containing $\log c$ and bounded by the (real) orthogonal complement of $\log(\overline{cp}) = \overline{\log c, 0}$.

The right vertical arrow is a weak homotopy equivalence by a standard scaling argument: any compact set in T_pX can be scaled down until it lies in $\log U$, and scaling preserves $\log \mathcal{H}_p$.

The same argument works for the left one: since the boundary of $\log B$ is tangent to the boundary of $\frac{1}{2}T_pX$ at 0, it is easy to see that any compact subset of $\frac{1}{2}T_pX$ can be sent into $\log B \cap \log U$ by multiplying it by a sufficiently small scalar. (One might worry about basic properties of ∂B like smoothness, since B was defined in terms of A, on which we made no smoothness assumptions. One can circumvent all worry by observing the following consequence of the triangle inequality: for small ε, B contains the open ε-ball D around the point of $\overline{Ap}$ at distance ε from p. And the boundary of $\log D$ is indeed smooth and tangent to $\frac{1}{2}T_pX$.)

Both weak homotopy equivalences are homotopy equivalences by Whitehead's theorem (or one could refine the scaling argument). To prove the theorem it now suffices to show that $\pi_1(T_pX - \log \mathcal{H}_p, \log c)$ is generated by the image of $\pi_1(\frac{1}{2}T_pX - \log \mathcal{H}_p, \log c)$, supplemented in case (2) by $\log \lambda$. This is just lemma 2.3, completing the proof. □

Proof of theorem 2.2. Let R be the set of $r \in (0, \infty]$ for which the conclusion of the theorem holds. By our assumption that some group of isometries of X preserves $\mathcal{M}$ and A and acts cocompactly on ∂A, the distances $d(A, H)$ are bounded away from 0, as H varies over $\mathcal{M}$. Therefore $B_r \cap \mathcal{H} = \emptyset$ for all sufficiently small r. It follows that R contains all small enough r. We will show below that that if $r \in R - \{\infty\}$ then $[r, r + \delta) \subseteq R$ for some $\delta > 0$. Also, we obviously have $B_r = \cup_{q<r} B_q$ for any $r \in (0, \infty]$. Therefore $(0, r) \subseteq R$ implies $(0, r] \subseteq R$. The connectedness of $(0, \infty]$ then implies $R = (0, \infty]$, proving the theorem.

So fix $r \in R - \{\infty\}$; we will exhibit $\delta > 0$ such that $[r, r + \delta) \subseteq R$. By $r \in R$ we know that $\pi_1(B_r - \mathcal{H}, A)$ is generated by the $\overline{AH}^{\circ}$'s for which $d(A, H) < r$. We abbreviate B_r to B and define S as the "sphere" ∂B. For each $p \in S$ there is an open ball U_p centered at p such that $U_p \cap \mathcal{H} = U_p \cap \mathcal{H}_p$. The cocompact action on ∂A we used in the previous paragraph is also cocompact on S. Also, we may choose the balls U_p so that the set of all of them is preserved by this action. It follows that there exists $\delta > 0$ such that $B_{r+\delta}$ is covered by B and all the U_p's.

To prove $[r, r+\delta) \subseteq R$ we suppose given some $r' \in (r, r+\delta)$ and write B' for $B_{r'}$. Since B' is covered by B and the U_p's, every mirror that meets B' either meets B or is tangent to S. So we must prove that $\pi_1(B' - \mathcal{H}, A)$ is generated by $\pi_1(B - \mathcal{H}, A)$ and the $\overline{AH}^{\circ}$'s with $H \in \mathcal{M}$ tangent to S. Lemma 2.4 says that $\pi_1\big((B \cup U_p) - \mathcal{H}, A\big)$

is generated by $\pi_1(B-\mathcal{H}, A)$, supplemented by $\overline{AH}^{\circ}$ if p is the point of tangency of S with some $H \in \mathcal{M}$.

For $p \in S$ we define $V_p = (B \cup U_p) \cap B'$. It is easy to see that the inclusion $V_p - \mathcal{H} \to (B \cup U_p) - \mathcal{H}$ is a homotopy equivalence. (Retract points of $U_p - B'$ along geodesics toward p.) So $\pi_1(V_p - \mathcal{H}, A)$ is generated by $\pi_1(B-\mathcal{H}, A)$, supplemented by $\overline{AH}^{\circ}$ if p is the point of tangency of S with some $H \in \mathcal{M}$. Because $B' = \cup_{p\in S} V_p$, repeatedly using van Kampen's theorem shows that $\pi_1(B' - \mathcal{H}, A)$ is generated by the images therein of all the $\pi_1(V_p - \mathcal{H}, A)$, finishing the proof.

This use of van Kampen's theorem requires checking that every set got from the V_p's by repeated unions and intersections is connected. To help verify this, we call a subset Y of X star-shaped (around A) if it contains A and the geodesics $\overline{Ay}$ for all $y \in Y$. The lemma below shows that each $B \cup U_p$ is star-shaped. Intersecting with B' preserves star-shapedness and yields V_p. Since unions and intersections of star-shaped sets are again star-shaped, our repeated application of van Kampen's theorem is legitimate. □

Lemma 2.5. *In the notation of the previous proof, $B \cup U_p$ is star-shaped around A.*

Proof. We must show that $y \in U_p$ implies $\overline{Ay} \subseteq B \cup U_p$. It suffices to prove $\overline{yz} \subseteq U_p$, where z is the point of ∂B closest to y. (We remarked above that the convexity of A implies that of B, and the uniqueness of z then follows from [BH, Prop. II.2.4].) Note that $\overline{zp}$ lies in the closure of B, since z and p do and B (hence its closure) is convex.

Consider the triangle p, y, z in X and a comparison triangle, meaning a triangle p', y', z' in $\mathbb{R}^2$ with the same edge lengths. We write θ for the angle between $\overline{zy}$ and $\overline{zp}$ at z, and similarly for θ'. Since X is a CAT(0) metric space we have $\theta' \geq \theta$ by [BH, Prop. II.3.1]. And we have $\theta \geq \pi/2$, because otherwise we could shorten $\overline{yz}$ by moving z toward p. Therefore θ' is the largest angle of the comparison triangle, so $\overline{p'y'}$ is its longest edge. Since the two triangles have the same edge lengths, $\overline{pz}$ is shorter than $\overline{py}$, so $z \in U_p$. Then U_p's convexity shows that all of $\overline{yz}$ lies in U_p. □

3. Loops in quotients of arrangement complements

We continue using the notation X, $\mathcal{M}$ and $\mathcal{H}$ from the previous section. We also suppose a group $P\Gamma$ acts isometrically and properly discontinuously on X, preserving $\mathcal{H}$. At this point we have no group Γ in mind; the notation $P\Gamma$ is just for compatibility with sections 4–5. Our goal is to understand the orbifold fundamental group of $(X - \mathcal{H})/P\Gamma$. We use the following definition from [L3] and [Ba2]; more general formulations exist [R][Ka].

Fixing a basepoint $a \in X - \mathcal{H}$, consider the set of pairs (γ, g) where $g \in P\Gamma$ and γ is a path in $X - \mathcal{H}$ from a to $g(a)$. We regard one such pair as equivalent to another one (γ', g') if $g = g'$ and γ and γ' are homotopic in $X - \mathcal{H}$, rel endpoints. The orbifold fundamental group $G_a := \pi_1^{\text{orb}}\big((X - \mathcal{H})/P\Gamma, a\big)$ means the set of equivalence classes. The group operation is $(\gamma, g) \cdot (\gamma', g') = (\gamma \text{ followed by } g \circ \gamma', gg')$. Projection of (γ, g) to g defines a homomorphism $G_a \to P\Gamma$. It is surjective because $X - \mathcal{H}$ is connected. The kernel is obviously $\pi_1(X - \mathcal{H}, a)$, yielding the exact sequence

$$1 \to \pi_1(X - \mathcal{H}, a) \to G_a \to P\Gamma \to 1 \tag{3.1}$$

Although we don't need it, we remark that if a has trivial $P\Gamma$ stabilizer then there is a simpler $P\Gamma$-invariant description of the orbifold fundamental group. Writing o for a's orbit, we define $G_o := \pi_1^{\text{orb}}\big((X-\mathcal{H})/P\Gamma, o\big)$ as the set of $P\Gamma$-orbits on the homotopy classes (rel endpoints) of paths in $X-\mathcal{H}$ that begin and end in o. The $P\Gamma$-action is the obvious one: $g \in P\Gamma$ sends a path γ to $g\circ\gamma$. To define $\gamma\gamma'$, where $\gamma, \gamma' \in G_o$, one translates γ' so that it begins where γ ends, and then composes paths in the usual way. Well-definedness of multiplication, and the identification with the definition of G_a, uses the fact that every path starting in o has a unique translate starting at a.

A complex reflection means a finite-order isometry of X whose fixed-point set is a complex hyperplane, called its mirror. In our applications, $P\Gamma$ is generated by complex reflections whose mirrors are hyperplanes in $\mathcal{M}$. This leads to certain natural elements of the orbifold fundamental group: for $H \in \mathcal{M}$ we next define a loop $\mu_{a,H} \in G_a$ which is a fractional power of $\overline{a\overset{\circ}{H}}$. (The meridians of conjecture 1.3 and theorems 1.4–1.5 are special cases of these loops.) Write p for the point of H closest to a and n_H for the order of the cyclic group generated by the complex reflections in $P\Gamma$ with mirror H. Write R_H for the isometry of X that fixes H pointwise and acts on its normal bundle by $\exp(2\pi i/n_H)$. This is a complex reflection and lies in $P\Gamma$, except when H is not the mirror of any reflection in $P\Gamma$. In that case, R_H is the identity map.

Recall that the definition of $\overline{a\overset{\circ}{H}}$ involved a point c of $\overline{ap}$ very near p, and a circular loop in $\overline{ap}^{\mathbb{C}}$ centered at p and based at c. We define $\mu_{a,H}$ as $\widehat{ac}$ followed by the first $(1/n_H)$th of this loop (going from c to $R_H(c)$), followed by $R_H(\text{reverse}(\widehat{ac}))$. (One can see such a path in figure 3.1, although the notation there is intended for a more elaborate situation considered below. The portion of the path in the figure that goes from b to $R_H(b)$ is $\mu_{b,H}$.) This is a path from a to $R_H(a)$, so the pair $(\mu_{a,H}, R_H)$ is an element of the orbifold fundamental group G_a. Using the definition of multiplication, the first component of $(\mu_{a,H}, R_H)^{n_H}$ is the path got by following $\mu_{a,H}$, then $R_H(\mu_{a,H})$, then $R_H^2(\mu_{a,H}), \ldots$ and finally $R_H^{n_H-1}(\mu_{a,H})$. It is easy to see that this is homotopic to $\overline{a\overset{\circ}{H}}$. So we have $(\mu_{a,H}, R_H)^{n_H} = \overline{a\overset{\circ}{H}}$.

At this point we have defined everything in the statement of theorem 1.2. But before proving it, we will adapt our construction to accomodate the "fat basepoints" of the previous section. This is necessary for our application to $\mathbb{C}H^{13}$. So we fix A as in section 2, and assume it contains our basepoint a. We will use A as the base "point" when discussing $\pi_1(X-\mathcal{H})$, and a as the basepoint when discussing $\pi_1^{\text{orb}}\big((X-\mathcal{H})/P\Gamma\big)$. In particular, the left term of (3.1) could also be written $\pi_1(X-\mathcal{H}, A)$. The analogue of $\mu_{a,H}$ is defined as follows, in terms of the point b of A that is closest to H. We define $\mu_{a,A,H}$ to be $\overline{ab}$ followed by $\mu_{b,H}$ followed by $R_H(\overline{ba})$. See figure 3.1 for a picture. Essentially the same argument as before shows that $(\mu_{a,A,H}, R_H)^{n_H} = \overline{A\overset{\circ}{H}} \in \pi_1(X-\mathcal{H}, A)$.

In applications one typically has some distinguished set of $\mu_{a,H}$'s or $\mu_{a,A,H}$'s in mind and wants to prove that they generate G_a. Theorem 1.2 in the introduction is a result of this sort, and the rest of the section is devoted to proving it. The following lemma is really the inductive step in the proof, so the reader might prefer to read the theorem's proof first. Also, theorem 1.2 uses only case (1) of the lemma; the other cases are for our application to $\mathbb{C}H^{13}$ in section 5.

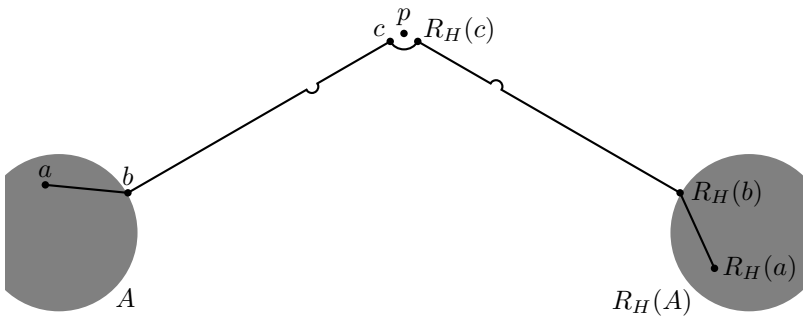

FIGURE 3.1. The path $\mu_{a,A,H}$ goes from left to right. Here R_H is the complex reflection of order 3, acting by counter-clockwise rotation by $2\pi/3$. The hyperplane H is not shown except for its point p closest to A. The small semicircles indicate that the path from b to c may detour around some points of $\mathcal{H}$.

Lemma 3.1. *Suppose $\mathcal{C} \subseteq \mathcal{M}$ are the hyperplanes closest to A, and let G be the subgroup of $G_a = \pi_1^{\text{orb}}\big((X-\mathcal{H})/P\Gamma, a\big)$ generated by the $(\mu_{a,A,C}, R_C)$ with $C \in \mathcal{C}$. Suppose $H \in \mathcal{M}$, write p for the closest point of H to A, r for $d(A,p)$, and B for the open r-neighborhood of A. Suppose G contains $\pi_1(B-\mathcal{H}, A)$ and that there exists a complex reflection $R \in P\Gamma$ with mirror in $\mathcal{C}$, such that one of the following holds:*

1. *R moves A closer to p.*
2. *R moves A closer to H, and no farther from p.*
3. *There exists an open ball U around p such that $U \cap \mathcal{H} = U \cap \mathcal{H}_p$, $B \cap R(B) \cap U \neq \emptyset$, and $R(B) \cap U \cap H \neq \emptyset$.*

Then G contains $\overline{AH}^{\circ}$.

Proof. In every case we have $R(B) \cap H \neq \emptyset$, so $R^{-1}(H)$ is closer to A than H is. Therefore $H \notin \mathcal{C}$, or in other words: the hyperplanes in $\mathcal{C}$ lie at distance $< r$ from A. We will prove the lemma under hypothesis (3), and then show that the other two cases follow. We hope figure 3.2 helps the reader. First we introduce the various objects pictured. As we did above, we write b for the point of A closest to H. Under our identification of $\pi_1(X-\mathcal{H}, A)$ with $\pi_1(X-\mathcal{H}, a)$, the loop $\overline{AH}^{\circ}$ corresponds to $\overline{ab}$ followed by $\overline{bH}^{\circ}$ followed by $\overline{ba}$.

The complex reflection R equals R_C^i for some $C \in \mathcal{C}$. The point marked C in the figure represents the point of C nearest to A. It lies inside B by the previous paragraph's remark that the elements of $\mathcal{C}$ are closer to A than H is. It also lies in $R(B)$, since R fixes C pointwise.

Consider the first component of $(\mu_{a,A,C}, R_C)^i$. After a homotopy it may be regarded as a path μ_1 in $B-\mathcal{H}$ from a to a point $c \in \big(B \cap R(B)\big) - \mathcal{H}$, followed by a path μ_2 in $R(B)-\mathcal{H}$ from c to $R(a)$. These paths are marked in the figure. So $(\mu_{a,A,C}, R_C)^i = (\mu_1\mu_2, R)$ in G_a.

The hypothesis that $B \cap R(B) \cap U \neq \emptyset$ is exactly what we need to know that some point y lies in this intersection, as drawn. The connectedness of $U \cap R(B)$ and the hypothesis that H meets $U \cap R(B)$ are exactly what we need to construct

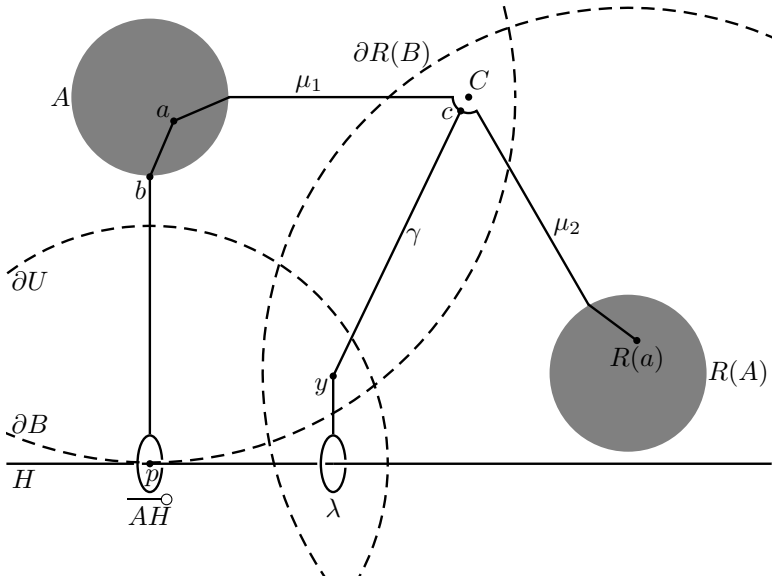

FIGURE 3.2. Illustration for the proof of lemma 3.1.

a loop λ in $(R(B) \cap U) - \mathcal{H}$, based at y, with linking number 1 with H. Finally, the connectedness of $B \cap R(B)$ allows us to construct a path γ in $(B \cap R(B)) - \mathcal{H}$ from c to y. This finishes the construction of the objects in the figure.

Our goal is to prove that G contains $\overline{AH}^{\circ}$. Lemma 2.4 shows that this loop lies in the subgroup of $\pi_1(X - \mathcal{H}, a)$ generated by $\pi_1(B - \mathcal{H}, a)$ and $\mu_1\gamma\lambda\gamma^{-1}\mu_1^{-1}$. This uses our hypothesis $U \cap \mathcal{H} = U \cap \mathcal{H}_p$. Since we assumed G contains the image of $\pi_1(B - \mathcal{H}, a)$, it suffices to show that G contains $\mu_1\gamma\lambda\gamma^{-1}\mu_1^{-1}$, or equivalently the homotopic loop $(\mu_1\mu_2)\big(\mu_2^{-1}\gamma\lambda\gamma^{-1}\mu_2\big)(\mu_2^{-1}\mu_1^{-1})$.

An element of the orbifold fundamental group G_a is really a pair, so we must prove $\big((\mu_1\mu_2)(\mu_2^{-1}\gamma\lambda\gamma^{-1}\mu_2)(\mu_2^{-1}\mu_1^{-1}), 1\big) \in G$. One checks that this equals

$$(\mu_1\mu_2, R) \cdot \Big(R^{-1}\big(\mu_2^{-1}\gamma\lambda\gamma^{-1}\mu_2\big), 1\Big) \cdot \Big(R^{-1}\big(\mu_2^{-1}\mu_1^{-1}\big), R^{-1}\Big)$$

The last term is the inverse of the first, which G contains by definition. So it suffices to show that the middle term lies in G, which is easy: the loop $\mu_2^{-1}\gamma\lambda\gamma^{-1}\mu_2$ lies in $R(B) - \mathcal{H}$, so its image under R^{-1} lies in $B - \mathcal{H}$. This finishes case (3).

Next we claim that (1) implies (3). Take U to be any ball around p with $U \cap \mathcal{H} = U \cap \mathcal{H}_p$. Then the remaining hypotheses of (3) follow immediately from $p \in R(B)$.

Finally we claim that (2) implies (3). By the previous paragraph it suffices to treat the case that $p \in \partial R(B)$. Take U to be any ball around p with $U \cap \mathcal{H} = U \cap \mathcal{H}_p$. The hypothesis $d(R(A), H) < r$ says that H is not orthogonal to $\overline{R(A), p}$. It follows that $R(B)$ contains elements of H arbitrarily close to p, so $U \cap R(B) \cap H \neq \emptyset$. Similarly, $d(R(A), H) < r$ implies the non-tangency of ∂B and $\partial R(B)$ at p. From this it follows that $B \cap R(B)$ has elements arbitrarily close to p, hence in U. This finishes the proof. □

Proof of theorem 1.2. We will apply lemma 3.1 with $A = \{a\}$, noting that $\mu_{a,A,H} = \mu_{a,H}$ for all $H \in \mathcal{M}$. Write G for the subgroup of G_a generated by the $(\mu_{a,C}, R_C)$'s.

By the exact sequence (3.1) and the assumed surjectivity $G \to P\Gamma$, it suffices to show that G contains $\pi_1(X - \mathcal{H}, a)$. By theorem 1.1 it suffices to show that it contains every $\overline{aH}$. We do this by induction on $d(a, H)$.

The base case is $H \in \mathcal{C}$, for which we use the fact that $\overline{aH}$ is a power of $(\mu_{a,H}, R_H)$. So suppose $H \in \mathcal{M} - \mathcal{C}$ and set $r := d(a, H)$. We may assume, by theorem 2.2 and the inductive hypothesis, that G contains $\pi_1(B - \mathcal{H}, a)$, where B is the open r-neighborhood of a. Then case (1) of lemma 3.1 shows that G also contains $\overline{aH}$, completing the inductive step. □

4. A monstrous(?) hyperplane arrangement

In this section we give background information on the conjecturally-monstrous hyperplane arrangement in $\mathbb{C}H^{13}$ which is the subject of conjecture 1.3 and theorems 1.4 and 1.5. For more information, see [A4][Ba1][Ba2][A5][H][HR].

We write $\mathbb{C}^{n,1}$ for a complex vector space equipped with a Hermitian form $\langle \cdot \,|\, \cdot \rangle$ of signature $(n, 1)$, assumed linear in its first argument and antilinear in its second. The *norm* v^2 of a vector v means $\langle v \,|\, v \rangle$. Complex hyperbolic space $\mathbb{C}H^n$ means the set of negative-definite 1-dimensional subspaces. If $V, W \in \mathbb{C}H^n$ are represented by vectors v, w then their hyperbolic distance is

$$d(V, W) = \cosh^{-1} \sqrt{\frac{\big|\langle v \,|\, w \rangle\big|^2}{v^2 w^2}} \tag{4.1}$$

If s is a vector of positive norm, then $s^\perp \subseteq \mathbb{C}^{n,1}$ defines a hyperplane in $\mathbb{C}H^n$, also written $s^\perp$, and

$$d(V, s^\perp) = \sinh^{-1} \sqrt{-\frac{\big|\langle v \,|\, s \rangle\big|^2}{v^2 s^2}} \tag{4.2}$$

These formulas are from [Go], up to an unimportant factor of 2.

A null vector means a nonzero vector of norm 0. If v is one then it represents a point V of the boundary $\partial\mathbb{C}H^n$. For any vector w of non-zero norm we define the *height* of w with respect to v by

$$\mathrm{ht}_v(w) := -\frac{\big|\langle v \,|\, w \rangle\big|^2}{w^2}. \tag{4.3}$$

This function is invariant under rescaling w, so it descends to a function on $\mathbb{C}H^n$, which is positive. The horosphere centered at V, of height h with respect to v, means the set of $p \in \mathbb{C}H^n$ with $\mathrm{ht}_v(p) = h$. We define open and closed horoballs the same way, replacing $=$ by $<$ and $\leq$. (More abstractly, one can define horospheres as the orbits of the unipotent radical of the $\mathrm{PU}(n, 1)$-stabilizer of V.)

We think of V as the center of these horospheres and horoballs and h as a sort of generalized radius, even though strictly speaking the distance from any point of $\mathbb{C}H^n$ to V is infinite. In particular, if $p, p' \in \mathbb{C}H^n$ then we say that p is closer to V than p' is, if $\mathrm{ht}_v(p) < \mathrm{ht}_v(p')$. To see that this notion depends on V rather than v, one checks that replacing v by a nonzero scalar multiple of itself does not affect this inequality. (It multiplies both sides by the same positive number.) Another way to think about this, at least for points outside some fixed closed horoball A centered at V, is to regard "closer to V" as alternate language for "closer to A".

In any case, in our application there will be a canonical choice for v, up to roots of unity.

Next we will describe the hyperplane arrangement appearing in conjecture 1.3 and theorems 1.4–1.5. We write ω for a primitive cube root of unity and define the Eisenstein integers $\mathcal{E}$ as $\mathbb{Z}[\omega]$. The Eisenstein integer $\omega - \bar{\omega} = \sqrt{-3}$ is so important that it has its own name θ. An *$\mathcal{E}$-lattice* means a free $\mathcal{E}$-module L equipped with a Hermitian form taking values in $\mathcal{E} \otimes \mathbb{Q} = \mathbb{Q}(\sqrt{-3})$, denoted $\langle \cdot \,|\, \cdot \rangle$. Sometimes we think of lattice elements as column vectors and $\langle \cdot \,|\, \cdot \rangle$ as specified by a matrix M equal to the transpose of its complex conjugate. Then $\langle v \,|\, w \rangle = v^T M \bar{w}$.

We will describe two $\mathcal{E}$-lattices, from [A5] and [A2]. Each has signature $(13, 1)$ and is equal to θ times its dual lattice. By [Ba1] there is only one lattice with these properties, so we may regard them as two different descriptions of the same lattice L. We will not actually use this uniqueness in this paper, and the first description of L is presented only to make precise the statement of theorem 1.4.

The definitions of $P\Gamma$ and $\mathcal{M}$, and some language we will use, are independent of the model. We regard $L \otimes_{\mathcal{E}} \mathbb{C}$ as a copy of $\mathbb{C}^{13,1}$, and take Γ to be the group of $\mathcal{E}$-linear automorphisms of L that preserve the inner product. As usual, $P\Gamma$ means the quotient by its subgroup of scalars. A *root* means a norm 3 lattice vector, the hyperplane arrangement $\mathcal{M}$ consists of the orthogonal complements in $\mathbb{C}H^{13}$ of the roots, and $\mathcal{H}$ means the union of these hyperplanes. The subject of conjecture 1.3 is the orbifold fundamental group of $(\mathbb{C}H^{13} - \mathcal{H})/P\Gamma$.

The special role of norm 3 vectors, and the name "root", arises as follows. First let $s \in \mathbb{C}^{13,1}$ be any vector of positive norm. Then the linear map

$$x \mapsto x - (1 - \omega)\frac{\langle x \,|\, s \rangle}{s^2} s$$

is an isometry of $\langle \,|\, \rangle$, called the *ω-reflection in s* and denoted R_s. Replacing ω by $\bar{\omega}$ gives the $\bar{\omega}$-reflection in s, which is the inverse of R_s. They are called triflections, because they are complex reflections of order 3. To see that R_s is a complex reflection (in particular an isometry) one checks that it fixes $s^\perp$ pointwise and multiplies s by ω.

In the special case that s is a root, R_s preserves L because of a conspiracy among the coefficients. First, the factor $(1-\omega)$ is a unit multiple of $\theta = \sqrt{-3}$. Second, for any lattice vector x, $\langle x \,|\, s \rangle$ is divisible by θ, since all inner products in L are. (This is what it means for L to lie in θ times its dual lattice.) Together these two factors of θ cancel the $s^2 = 3$ term in the denominator, up to a unit. So $R_s(x)$ is an $\mathcal{E}$-linear combination of x and s, hence lies in L. When one has a reflection in mind (real or complex), it is customary to call a vector orthogonal to its fixed-point set a root. When one also has a lattice in mind, one usually fixes the scale of a root by requiring it to be a primitive lattice vector. This is why we call norm 3 vectors roots. One can show that no other elements of $P\Gamma$ act on $\mathbb{C}H^{13}$ by complex reflections. (The analogous result for unimodular $\mathcal{E}$-lattices is contained in lemmas 8.1–8.2 of [A1]; for the current case one uses the fact that L is equal to θ times its dual lattice, rather than merely lying in it.) So $\mathcal{M}$ is exactly the set of mirrors of the complex reflections in $P\Gamma$, making $\pi_1^{\mathrm{orb}}\big((\mathbb{C}H^{13} - \mathcal{H})/P\Gamma\big)$ a braid-like group in the sense of this paper.

The second author showed in [Ba2, Lemma 3.3] that the triflections in a particular set of 14 roots generate Γ. The first author showed in [A2] that all roots are equivalent under Γ.

Our first description of L is the "$P^2\mathbb{F}_3$ model" from [A5]. As mentioned above, we include it only to give precise meaning to theorem 1.4, and we will not refer to it later. We start with the diagonal inner product matrix $[-1;1,\dots,1]$ on $\mathbb{C}^{13,1}$, and regard the last 13 coordinates as being indexed by the 13 points of $P^2\mathbb{F}_3$. The "point roots" are the vectors of the form $(0;\theta,0^{12})$ with the θ in any of last 13 positions. The "line roots" are the vectors of the form $(1;1,1,1,1,0,\dots,0)$, with 1's in positions corresponding to the points of a line in $P^2\mathbb{F}_3$. L is defined as the span of these 26 roots. This construction obviously has $\mathrm{PGL}_3(\mathbb{F}_3)$ symmetry, and it also has less-obvious symmetries exchanging the point roots and line roots (up to scalars). This yields a subgroup $\mathrm{PGL}_3(\mathbb{F}_3)\rtimes\mathbb{Z}/2$ of $P\Gamma$ that acts transitively on these 26 roots. We take τ in theorem 1.4 to be the unique point of $\mathbb{C}H^{13}$ invariant under this group; its coordinates are $(4+\sqrt{3};1,\dots,1)$. Among other results, it was shown in [Ba1], Prop. 1.2, that the hyperplanes in $\mathcal{M}$ that are closest to τ are exactly the mirrors of the point and line roots. We have now described concretely all the objects in theorem 1.4.

Now we give the "Leech model" of L from [A2], and make concrete the objects in theorem 1.5. We will use this model for the rest of the paper. We define L as the $\mathcal{E}$-lattice $\Lambda\oplus\left(\begin{smallmatrix}0&\bar\theta\\ \theta&0\end{smallmatrix}\right)$, where Λ is the complex Leech lattice at the smallest scale at which all inner products lie in $\mathcal{E}$. The complex Leech lattice is studied in detail in [W]; the scale used there is the most natural one for coordinate computations in it, and has minimal norm 9. At our scale it has minimal norm 6, all inner products are divisible by θ, and Λ is equal to θ times its dual lattice. It is called the complex Leech lattice because its underlying real lattice is a scaled copy of the usual (real) Leech lattice described in [C]. The properties of Λ that we will use are that its automorphism group is transitive on its vectors of norms 6 and 9, and that its covering radius is $\sqrt{3}$. The transitivity is proven in [W]. The meaning of the covering radius is that closed balls of that radius, centered at lattice points, exactly cover Euclidean space. It is $\sqrt{3}$ because the real Leech lattice, at its own natural scale, has minimal norm 4 and (by [CPS]) covering radius $\sqrt{2}$.

We will write vectors of L in the form $(x;y,z)$, where $x\in\Lambda$ and $y,z\in\mathcal{E}$.

Conceptually, our "basepoint" for the description of generators for $\pi_1^{\mathrm{orb}}\big((\mathbb{C}H^{13}-\mathcal{H})/P\Gamma\big)$ in theorem 1.5 is the cusp of $P\Gamma$ represented by the null vector $\rho:=(0;0,1)$. As explained in the introduction, really this is a shorthand for choosing a "fat basepoint": a closed horoball A centered at ρ that misses $\mathcal{H}$. The following lemma shows that such a horoball exists. More precisely, it identifies the largest open horoball centered at ρ that misses $\mathcal{H}$; we may take A to be any closed horoball inside it. We also fix a basepoint $a\in A$, so now the paths $\mu_{a,A,H}$ in theorem 1.5 have been defined.

Lemma 4.1. *The open horoball $\big\{p\in\mathbb{C}H^{13}\;\big|\;\mathrm{ht}_\rho(p)<1\big\}$ is disjoint from $\mathcal{H}$, and the mirrors that meet its boundary are the orthogonal complements of the roots l that satisfy $\big|\langle\rho\,|\,l\rangle\big|^2=3$.*

Proof. The special property of ρ we need is that it is orthogonal to no roots. This is clear because $\rho^\perp=\Lambda\oplus\langle\rho\rangle$ has no vectors of norm 3. Now, if l is a root then the point of l's mirror nearest to ρ is represented by the vector projection of ρ to $l^\perp$, namely $p=\rho-\frac{1}{3}\langle\rho\,|\,l\rangle l$. One computes $\mathrm{ht}_\rho(p)=\big|\langle\rho\,|\,l\rangle\big|^2/3$. This is at least 1, with equality just if $\big|\langle\rho\,|\,l\rangle\big|^2=3$. □

Lemma 4.1 also identifies the mirrors closest to ρ, namely the orthogonal complements of the roots l with $\langle \rho \,|\, l \rangle$ equal to a unit multiple of θ. After scaling l we may suppose

$$l = \left(\lambda; 1, \theta\Big(\frac{\lambda^2 - 3}{6} + \nu_l\Big)\right) \tag{4.4}$$

where $\lambda \in \Lambda$ and ν_l is purely imaginary and chosen so that the last coordinate lies in $\mathcal{E}$. The set of possibilities for ν_l is $\frac{1}{\theta}(\frac{1}{2} + \mathbb{Z})$ if 6 divides λ^2 and $\frac{1}{\theta}\mathbb{Z}$ otherwise. Despite its elaborate form, the last coordinate is well-suited for the calculations required in next section. We call these roots the *Leech roots* (hence the notation l), their mirrors the *Leech mirrors* and the meridians $(\mu_{a,A,l^\perp}, R_l)$ the *Leech meridians*. We have now made concrete all the objects in theorem 1.5.

We remark that there is a 25-dimensional integral Heisenberg group in $P\Gamma$ that acts simply-transitively on the Leech roots. It consists of the "translations" in the proof of lemma 5.1. Conceptually, this is a simpler way to index the Leech roots than by the pairs λ, ν_l, but in the end it is equivalent. Also, these translations act cocompactly on ∂A, verifying the technical condition we required on A in section 2.

Early in the section we explained how one can meaningfully say that one point of $\mathbb{C}H^{13}$ is closer to a point of $\partial\mathbb{C}H^{13}$ than another point of $\mathbb{C}H^{13}$ is. We will also need to be able to compare the "distance" from a point $p \in \mathbb{C}H^{13}$ to two different cusps $V, V' \in \partial\mathbb{C}H^{13}$ of $P\Gamma$. Being cusps, they can be represented by lattice vectors v, v', which we may choose to be primitive. Then v, v' are well-defined up to multiplication by sixth roots of unity, and the corresponding height functions ht_v, $\mathrm{ht}_{v'}$ are independent of these factors. So we will say that p is closer to V than to V' if $\mathrm{ht}_v(p) < \mathrm{ht}_{v'}(p)$. Note that this construction depends on the fact that V, V' are cusps; it does not make sense for general points of $\partial\mathbb{C}H^{13}$. In our applications, v and v' will always be ρ and a translate of ρ.

5. The Leech meridians generate

The purpose of this section is to prove theorem 1.5, showing that the Leech meridians generate the orbifold fundamental group $G_a := \pi_1^{\mathrm{orb}}\big((\mathbb{C}H^{13} - \mathcal{H})/P\Gamma, a\big)$, where $\mathcal{M}$, $\mathcal{H}$, $P\Gamma$, A, a and the Leech meridians $(\mu_{a,A,l^\perp}, R_l)$ are defined in the previous section. We write G for the subgroup of G_a they generate, and we must prove that G is all of G_a.

We begin with an overview of the proof, which follows that of theorem 1.2. It amounts to showing that the mirror of any non-Leech root s satisfies one of the hypotheses (1)–(3) of lemma 3.1. It turns out (lemma 5.2) that if $\big|\langle \rho \,|\, s \rangle\big|^2 > 21$ then the simplest hypothesis (1) holds. If $\big|\langle \rho \,|\, s \rangle\big|^2 > 9$ then the same method shows that the next simplest hypothesis (2) holds (lemmas 5.2 and 5.3). For the case $\big|\langle \rho \,|\, s \rangle\big|^2 = 9$ we enumerate the orbits of roots $(?; \theta, ?)$ under the Γ-stabilizer of ρ (lemma 5.1). There are three orbits, satisfying hypotheses (1), (2) and (3) of lemma 3.1, respectively. The last orbit is especially troublesome (lemma 5.4). The proof of theorem 1.5 is then a wrapper around these results.

The following description of vectors in $L \otimes \mathbb{C}$ is very important in our computations. We will use it constantly, often specializing to the case of roots. Every

vector $s \in (L \otimes \mathbb{C}) - \rho^\perp$ can be written uniquely in the form

$$s = \Big(\sigma; m, \frac{\theta}{\bar m}\Big(\frac{\sigma^2 - N}{6} + \nu\Big)\Big) \tag{5.1}$$

where $\sigma \in \Lambda \otimes \mathbb{C}$, $m \in \mathbb{C} - \{0\}$, N is the norm s^2, and ν is purely imaginary. Restricting the first coordinate to Λ and the others to $\mathcal{E}$ gives the elements of $L - \rho^\perp$. Further restricting N to 3 gives the roots of L, and finally restricting m to 1 gives the Leech roots from (4.4). For vectors of any fixed negative (resp. positive) norm, the larger the absolute value of the middle coordinate m, the further from ρ lie the corresponding points (resp. hyperplanes) in $\mathbb{C}H^{13}$.

One should think of s from (5.1) as being associated to the vector σ/m in the positive-definite Hermitian vector space $\Lambda \otimes_{\mathcal{E}} \mathbb{C}$. By this we mean that the most important part of $\langle s \,|\, s'\rangle$ is governed by the relative positions of σ/m and σ'/m'. Namely, by writing out $\langle s \,|\, s'\rangle$, completing the square and patiently rearranging, one can check

$$\langle s \,|\, s'\rangle = m\bar m' \Big[\frac{1}{2}\Big(\frac{N'}{|m'|^2} + \frac{N}{|m|^2} - \Big(\frac{\sigma}{m} - \frac{\sigma'}{m'}\Big)^2\Big) + \operatorname{Im}\Big\langle \frac{\sigma}{m} \;\Big|\; \frac{\sigma'}{m'}\Big\rangle + 3\Big(\frac{\nu'}{|m'|^2} - \frac{\nu}{|m|^2}\Big)\Big]. \tag{5.2}$$

Caution: we are using the convention that the imaginary part of a complex number is imaginary; for example $\operatorname{Im}\theta$ is θ rather than $\sqrt{3}$.

In the rest of this section, "s" will only be used for roots.

Lemma 5.1. *Suppose λ_6, λ_9 are fixed vectors in Λ with norms 6 and 9. Then under the Γ-stabilizer of ρ, every root with $m = \theta$ is equivalent to $(0; \theta, -\omega)$ or $(\lambda_6; \theta, \omega)$ or $(\lambda_9; \theta, -1)$.*

Proof. The Γ-stabilizer of ρ contains the Heisenberg group of "translations"

$$\begin{aligned} (l; 0, 0) &\mapsto \big(l; 0, \bar\theta^{-1}\langle l \,|\, \lambda\rangle\big) \\ T_{\lambda,z} : (0; 1, 0) &\mapsto \big(\lambda; 1, \theta^{-1}(z - \lambda^2/2)\big) \\ (0; 0, 1) &\mapsto (0; 0, 1) \end{aligned}$$

where $\lambda \in \Lambda$ and $z \in \operatorname{Im}\mathbb{C}$ are such that $z - \lambda^2/2 \in \theta\mathcal{E}$. Suppose $s \in L$ has the form (5.1) with $N = 3$ and $m = \theta$. Applying $T_{\lambda,z}$ to s changes the first coordinate by $\theta\lambda$. By [W, p. 153], every element of Λ is congruent modulo $\theta\Lambda$ to a vector of norm 0, 6 or 9, so we may suppose σ has one of these norms. Since $\operatorname{Aut}\Lambda$ fixes ρ and acts transitively on the vectors of each of these norms [W, p. 155], we may suppose $s = 0$, λ_6 or λ_9. That is, s is one of

$$(0; \theta, \tfrac{1}{2} - \nu) \qquad (\lambda_6; \theta, -\tfrac{1}{2} - \nu) \qquad (\lambda_9; \theta, -1 - \nu).$$

In each of the three cases, the possibilities for ν differ by the elements of $\operatorname{Im}\mathcal{E}$. Applying $T_{0,z}$ ($z \in \operatorname{Im}\mathcal{E}$) adds z to the third coordinate of s. Therefore we may take $\nu = \theta/2$, $\bar\theta/2$ and 0 in the three cases, yielding the roots in the statement of the lemma. (These three roots are inequivalent under the Γ-stabilizer of ρ, but we don't need this.) □

Lemma 5.2. *Suppose s is the root $(0; \theta, -\omega)$ or a root as in 5.1 with $|m| = 2$ or $|m| > \sqrt{7}$, and define p as the point of $s^\perp$ nearest ρ. Then there is a triflection in a Leech root that moves ρ closer to p.*

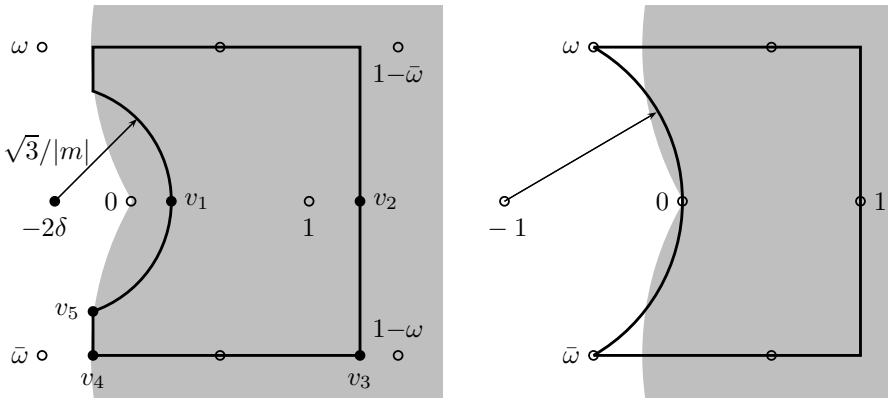

FIGURE 5.1. See the proof of lemma 5.2. V is the union of the gray (open) disks, which have radius $\sqrt{3}$ and centers $1-\omega^{\pm 1}$. We seek a Leech root l so that y lies in this region. U is the closed region bounded by the solid line, and is where we can arrange for y to be. U varies with $|m|$; we have drawn the case $|m| = \sqrt{7}$, when v_5 is on the boundary of V, and the case $|m| = \sqrt{3}$, when v_4 and v_5 coalesce at $\bar{\omega}$. Hollow circles indicate Eisenstein integers.

Proof of lemma 5.2. This proof grew from simpler arguments used for [A2, Thm. 4.1] and [Ba1, Prop. 4.2].

We have $p = \rho - \frac{1}{3}\langle \rho \,|\, s\rangle s = \rho + \frac{1}{\theta}\bar{m}s$. We want to choose a Leech root l, and $\zeta = \omega^{\pm 1}$, such that the ζ-reflection in l (call it R) moves ρ closer to p. This is equivalent to $\langle p \,|\, R(\rho)\rangle$ being smaller in absolute value than $\langle p \,|\, \rho\rangle$. We will write down these inner products explicitly and then choose l and ζ appropriately. Direct calculation gives $\langle p \,|\, \rho\rangle = -|m|^2$. Also,

$$R(\rho) = \rho - (1-\zeta)\frac{\langle \rho \,|\, l\rangle}{\langle l \,|\, l\rangle} l = \rho + \frac{1-\zeta}{\theta} l.$$

It turns out that the necessary estimates on $\langle p \,|\, R(\rho)\rangle$ are best expressed in terms of the following parameter:

$$y := \frac{\theta}{|m|^2}\langle p \,|\, l\rangle = \frac{\theta}{|m|^2}\Big\langle \rho + \frac{\bar{m}}{\theta}s \;\Big|\; l \Big\rangle = -\frac{3}{|m|^2} + \frac{1}{m}\langle s \,|\, l\rangle \tag{5.3}$$

$$\in\; -\frac{3}{|m|^2} + \frac{1}{m}\theta\mathcal{E}. \tag{5.4}$$

First one works out

$$\left|\frac{\langle p \,|\, R(\rho)\rangle}{\langle p \,|\, \rho\rangle}\right| = \left|\frac{\langle p \,|\, R(\rho)\rangle}{|m|^2}\right| = \left|\tfrac{1}{3}(1-\bar{\zeta})y - 1\right|. \tag{5.5}$$

Our goal is to choose l and ζ so that this is less than 1. This is equivalent to $|y-(1-\zeta)| < \sqrt{3}$. Because the possibilities for ζ are $\omega^{\pm 1}$, this amounts to being able to choose l so that y lies in the union V of the open balls in $\mathbb{C}$ of radius $\sqrt{3}$ around the points $1-\omega$ and $1-\bar{\omega}$. So our goal is to choose l such that y lies in the shaded region in figure 5.1.

Now we examine how our choice of l affects y. Writing l as in (4.4), choosing it amounts to choosing $\lambda \in \Lambda$, and then choosing $\nu_l \in \operatorname{Im}\mathbb{C}$ subject to the condition

that the last coordinate of (4.4) is in $\mathcal{E}$. Specializing (5.2) to the case that s has norm $N = 3$ and s' is the Leech root l gives

$$\langle s \,|\, l\rangle = m\Big[\frac{3}{2|m|^2} + \frac{3}{2} - \frac{1}{2}\Big(\frac{\sigma}{m} - \lambda\Big)^2 + \mathrm{Im}\Big\langle \frac{\sigma}{m} \;\Big|\; \lambda\Big\rangle + 3\Big(\nu_l - \frac{\nu}{|m|^2}\Big)\Big].$$

Plugging this into formula (5.3) gives

$$y = -\frac{3}{2|m|^2} + \frac{3}{2} - \frac{1}{2}\Big(\frac{\sigma}{m} - \lambda\Big)^2 + \mathrm{Im}\Big\langle \frac{\sigma}{m} \;\Big|\; \lambda\Big\rangle + 3\Big(\nu_l - \frac{\nu}{|m|^2}\Big). \tag{5.6}$$

The covering radius of a lattice in Euclidean space is defined as the smallest number such that the closed balls of that radius, centered at lattice points, cover Euclidean space. The covering radius of Λ is $\sqrt{3}$, because the underlying real lattice has norms equal to $3/2$ times those of the real Leech lattice, whose covering radius is $\sqrt{2}$ by [CPS]. Therefore we may take λ so that $0 \le (\sigma/m - \lambda)^2 \le 3$. It follows that the real part of (5.6) lies in $[-\delta, 3/2 - \delta]$ where $\delta := 3/2|m|^2$.

Next we choose ν_l. The only constraint on it is that the last component of $l = (\lambda; 1, ?)$ must lie in $\mathcal{E}$. As mentioned after (4.4), this amounts to $\nu_l \in \frac{1}{\theta}(\frac{1}{2} + \mathbb{Z})$ if λ^2 is divisible by 6, and $\nu_l \in \frac{1}{\theta}\mathbb{Z}$ otherwise. In either case, referring to (5.6) shows that changing our choice of ν_l allows us to change y by any rational integer multiple of θ. So we may take $\mathrm{Im}\, y \in [-\theta/2, \theta/2]$. After these choices we have

$$\mathrm{Re}\, y \in [-\delta, 3/2 - \delta] \quad \text{and} \quad \mathrm{Im}\, y \in [-\theta/2, \theta/2]. \tag{5.7}$$

Now we can derive additional information about y. We have $y \ne -2\delta$ since -2δ is not in the rectangle (5.7), and since $y \in -2\delta + \frac{\theta}{m}\mathcal{E}$ by (5.4), y lies at distance $\ge \sqrt{3}/|m|$ from -2δ. We define U as the closed rectangle (5.7) in $\mathbb{C}$ minus the open $(\sqrt{3}/|m|)$-disk around -2δ. We have shown that we may choose a Leech root l such that $y \in U$. We have indicated U in outline in figure 5.1; as $|m|$ increases, the rectangle moves to the right, the center of the removed disk approaches zero, and its radius approaches zero more slowly than the center does.

Now suppose $|m|^2 > 7$. We claim $U \subseteq V$. Assuming this for the moment, we may choose l such that y is in U, hence V, which allows us to choose $\zeta = \omega^{\pm 1}$ so that ζ-reflection in l moves ρ closer to p. This finishes the proof. To prove the claim it will suffice to show that the lower half of U lies in the open $\sqrt{3}$-ball around $1 - \omega$. Obviously it suffices to check this for the points marked $v_1, \ldots, v_5$ in figure 5.1. These are $v_1 = -2\delta + \sqrt{3}/|m|$, $v_2 = \frac{3}{2} - \delta$, $v_3 = \frac{3}{2} - \delta - \frac{1}{2}i\sqrt{3}$, $v_4 = -\delta - \frac{1}{2}i\sqrt{3}$ and

$$v_5 = -\delta - i\sqrt{\frac{3}{|m|^2} - \frac{9}{4|m|^4}}.$$

Using $|m|^2 > 7$, one can check that each of these lies at distance $< \sqrt{3}$ from $1 - \omega$. This finishes the proof of the $|m|^2 > 7$ case. (If $|m|^2 = 7$ then v_5 lies in the boundary of V. If $|m|^2 = 4$ then v_4 and v_5 are outside the boundary; see figure 5.2. If $|m|^2 = 3$ then $v_1 = 0$ is on the boundary and $v_4 = v_5 = \bar{\omega}$ is outside it; see the second part of figure 5.1.)

Next we treat the special case $s = (0; \theta, -\omega)$. Choosing $\lambda = 0$ gives $\mathrm{Re}\, y = 1$ by (5.6). Then choosing ν_l as above, so that $\mathrm{Im}\, y$ lies in $[-\theta/2, \theta/2]$, yields $y \in V$. So we can move ρ closer to p just as in the $|m|^2 > 7$ case.

Finally, we suppose $|m| = 2$; we may take $m = 2$ by multiplying s by a unit. Recall that once we proved that y lies in the rectangle (5.7), we could use (5.4) to show that y lies outside the open disk used in the definition of U. For $m = 2$ the

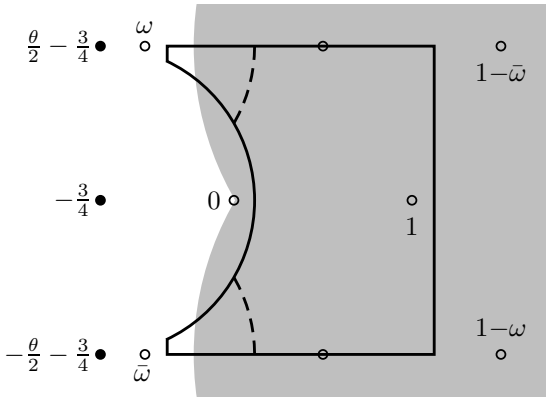

FIGURE 5.2. The analogue of figure 5.1 for the special case $|m| = 2$ in the proof of lemma 5.2. The proof shows that y lies in U (bounded by the solid path) but outside two open disks (indicated by the dashed arcs), hence in V (the shaded region).

argument shows more. Since $\delta = \frac{3}{8}$ when $|m| = 2$, (5.4) shows $y \in -\frac{3}{4} + \frac{\theta}{2}\mathcal{E}$. Since $-\frac{3}{4} \pm \frac{\theta}{2}$ lie in $-\frac{3}{4} + \frac{\theta}{2}\mathcal{E}$ but not in the rectangle (5.7), y lies at distance $\geq \sqrt{3}/2$ from each of them, just as it lies at distance $\geq \sqrt{3}/2$ from $-\frac{3}{4}$. It is easy to check that U minus the open $\sqrt{3}/2$-balls around $-\frac{3}{4} \pm \frac{\theta}{2}$ lies in V; see figure 5.2. Therefore $y \in V$, finishing the proof as before. (One can consider analogues of these extra disks for any m. They are unnecessary if $|m|^2 > 7$, and turn out to be useless if $|m|^2 = 3$ or 7.) □

Lemma 5.3. *Suppose s is the root $(\lambda_6; \theta, \omega)$ or a root as in 5.1 with $|m| = \sqrt{7}$, and define p as the point of $s^\perp$ nearest ρ. Then there is a triflection in a Leech root that either moves ρ closer to p, or else moves ρ closer to $s^\perp$ while preserving ρ's distance from p.*

Proof. Suppose first $|m| = \sqrt{7}$. Then the proof of lemma 5.2 goes through unless y is v_5 in figure 5.1, or its complex conjugate. So suppose $y = v_5$ or $\bar{v}_5$, and take $\zeta = \omega$ or $\bar{\omega}$ respectively. The argument in the proof of lemma 5.2, that R moves ρ closer to p, fails because $\big|y - (1-\zeta)\big|$ equals $\sqrt{3}$ rather than being strictly smaller. But it does show that $R(\rho)$ is exactly as far from p as ρ is. This is one of our claims, and what remains to show is that R moves ρ closer to $s^\perp$.

To do this we first solve (5.3) for $\langle s \,|\, l\rangle$ in terms of y, obtaining $\langle s \,|\, l\rangle = (3 + |m|^2 y)/\bar{m}$. Then one works out

$$\left|\frac{\langle s \,|\, R(\rho)\rangle}{\langle s \,|\, \rho\rangle}\right| = \left|\frac{\langle s \,|\, \rho\rangle - \frac{1}{\theta}(1-\bar{\zeta})\langle s \,|\, l\rangle}{\langle s \,|\, \rho\rangle}\right| = \left|1 - \frac{1}{3}(1-\bar{\zeta})\Big(\frac{3}{|m|^2} + y\Big)\right|.$$

We want this to be less than 1. By copying the argument following (5.5), this is equivalent to $y + 3/|m|^2$ lying in the open $\sqrt{3}$-disk around $1 - \zeta$. This is obvious from the figure because $y + 3/|m|^2$ is $3/7$ to the right of $y = v_5$ or $\bar{v}_5$. This finishes the $|m| = \sqrt{7}$ case.

The case $s = (\lambda_6; \theta, \omega)$ is similar. In this case U appears in figure 5.1. Writing l as in (4.4) with $\lambda = 0$ leads to $\operatorname{Re} y = 0$, so either $y \in V$ (so the proof of lemma 5.2

applies) or else $y = 0 \in \partial V$. In this case the argument for $|m| = \sqrt{7}$ shows that $R(\rho)$ is just as close to p as ρ is, and that R moves ρ closer to $s^\perp$. □

Lemma 5.4. *Let $s = (\lambda_9; \theta, -1)$, define p as the point of $s^\perp$ nearest ρ, and B as the open horoball centered at ρ, whose bounding horosphere is tangent to $s^\perp$ at p. Then there exists an open ball U around p with $U \cap \mathcal{H} = U \cap \mathcal{H}_p$, and a triflection R in one of the Leech mirrors, such that $B \cap R(B) \cap U \neq \emptyset$ and $R(B) \cap U \cap s^\perp \neq \emptyset$.*

Proof. Since we are verifying hypothesis (3) of lemma 3.1, we will use that lemma's notation H for $s^\perp$. By definition,

$$p = \rho - \tfrac{1}{3}\langle \rho \,|\, s\rangle s = (-\lambda_9; \bar{\theta}, 2).$$

This has norm -3 and lies in L. One computes $\mathrm{ht}_\rho(p) = 3$, so B is the height 3 open horoball around ρ. We take U to have radius $\sinh^{-1}\sqrt{1/3}$. To check that $U \cap \mathcal{H} = U \cap \mathcal{H}_p$, consider a root s' not orthogonal to p. Then $\big|\langle p \,|\, s'\rangle\big| \geq \sqrt{3}$ since $p \in L$, so

$$d(p, s'^\perp) = \sinh^{-1}\sqrt{-\frac{\big|\langle p \,|\, s'\rangle\big|^2}{p^2 s'^2}} \geq \sinh^{-1}\sqrt{1/3},$$

as desired.

Next we choose R to be the ω-reflection in the Leech root $l = (0; 1, -\omega)$. (We found l by applying the proof of lemma 5.2 as well as we could. That is, we choose l so that y in that proof equals the lower left corner $\bar{\omega}$ of the second part of figure 5.1.) This yields $R(\rho) = (0; \bar{\omega}, 0)$. We must verify $R(B) \cap U \cap H \neq \emptyset$ and $B \cap R(B) \cap U \neq \emptyset$.

Our strategy for $R(B) \cap U \cap H \neq \emptyset$ is to begin by defining p' as the projection of $R(\rho)$ to H, which turns out to lie outside U. Then we parameterize $\overline{p'p} \subseteq H$, find the point x where it crosses ∂U, and check that $x \in R(B)$. So $x \in R(B) \cap H \cap \partial U$. Therefore a point of $\overline{p'p}$, slightly closer to p than x is, lies in $R(B) \cap H \cap U$, showing that this intersection is nonempty.

Here are the details. Computation gives $p' = (\bar{\omega}\lambda_9/\theta; 2\bar{\omega}, -\bar{\omega}/\theta)$, of norm -1. One checks $\langle p' \,|\, p\rangle = 2\bar{\omega}\bar{\theta}$, so $d(p, p') = \cosh^{-1} 2 > \sinh^{-1}\sqrt{1/3}$ and p' lies outside U, as claimed. Also, $-\omega\theta p'$ and p have negative inner product. Therefore $\overline{p'p} - \{p\}$ is parameterized by $x_t = -\omega\theta p' + tp$ with $t \in [0, \infty)$. One computes $\langle x_t \,|\, p\rangle = -3t - 6$ and $x_t^2 = -3t^2 - 12t - 3$, yielding

$$d(x_t, p) = \cosh^{-1}\sqrt{\frac{\big|\langle x_t \,|\, p\rangle\big|^2}{x_t^2 p^2}} = \cosh^{-1}\sqrt{\frac{t^2 + 4t + 4}{t^2 + 4t + 1}}$$

Now, x_t lies in ∂U just when this equals $\sinh^{-1}\sqrt{1/3}$, yielding a quadratic equation for t. There is just one nonnegative solution, namely $t = 2\sqrt{3} - 2$. So $x = x_{2\sqrt{3}-2}$. Then one computes $\langle R(\rho) \,|\, x\rangle = \bar{\omega}\bar{\theta}(4\sqrt{3} - 3)$, so

$$\mathrm{ht}_{R(\rho)}(x) = -\frac{\big|\langle R(\rho) \,|\, x\rangle\big|^2}{x^2} = -\frac{3\big(57 - 24\sqrt{3}\big)}{-27} < 3.$$

That is, $x \in R(B)$ as desired.

Our strategy for $B \cap R(B) \cap U \neq \emptyset$ is similar. We parameterize the geodesic $\overline{x\rho}$, where x is the point found in the previous paragraph, find the point y where it crosses ∂B, and check that y lies in $R(B)$ and U. Here are the details. Computation shows $\langle x \,|\, \rho\rangle = -6\sqrt{3} < 0$, so $\overline{x\rho} - \{\rho\}$ is parameterized by $y_u = x + u\rho$ with $u \in [0, \infty)$. Further computation shows $\langle y_u \,|\, \rho\rangle = -6\sqrt{3}$ and $y_u^2 = -27 - 12u\sqrt{3}$,

so $\mathrm{ht}_\rho(y_u) = 36/(9+4u\sqrt{3})$. Setting this equal to 3 yields $u = \sqrt{3}/4$, so $y = y_{\sqrt{3}/4}$. Now one checks that $\mathrm{ht}_{R(\rho)}(y) < 3$, so that $y \in R(B)$. A similar calculation proves $y \in U$. (In fact this calculation can be omitted, because y, p are the projections to ∂B of the two points x, p outside B, but not both in ∂B. Projection to a closed horoball decreases the distance between two points, if at least one of them is outside it. Therefore $d(y,p) < d(x,p) = \sinh^{-1}\sqrt{1/3}$.) □

Proof of theorem 1.5. We will mimic the proof of theorem 1.2 (see the end of section 3), using lemmas 5.2–5.4 in place of the "moves a closer to p" hypothesis of that theorem. Write G for the subgroup of $G_a = \pi_1^{\mathrm{orb}}\big((\mathbb{C}H^{13} - \mathcal{H})/P\Gamma, a\big)$ generated by the Leech meridians, i.e., the pairs $(\mu_{a,A,l^\perp}, R_l)$ with l a Leech root. We must show that G is all of G_a. It is known ([Ba1], or [A5] for a later proof) that the R_l's generate $P\Gamma$. By the exact sequence (3.1), it therefore suffices to show that G contains $\pi_1(\mathbb{C}H^{13} - \mathcal{H}, a)$. By theorem 2.2 it suffices to show that G contains every $\overline{AH}^\circ$, with H varying over $\mathcal{M}$. We do this by induction on the distance from H to ρ, or properly speaking, on $\big|\langle \rho \,|\, s\rangle\big|$ where s is a root with $H = s^\perp$. The base case is when s is a Leech root, i.e., $\big|\langle \rho \,|\, s\rangle\big| = \sqrt{3}$, and we just observe $\overline{AH}^\circ = (\mu_{a,A,H}, R_s)^3$.

Now suppose s is a root but not a Leech root, $H = s^\perp$, p is the point of H closest to ρ, and B is the open horoball centered at ρ and tangent to H at p. We may assume by induction that G contains every $\overline{A, s'^\perp}^\circ$ with s' a root satisfying $\big|\langle \rho \,|\, s'\rangle\big| < \big|\langle \rho \,|\, s\rangle\big|$. It follows from theorem 2.2 that G contains $\pi_1(B - \mathcal{H}, a)$.

The smallest possible value of $\big|\langle \rho \,|\, s\rangle\big|$ for a non-Leech root s is 3, occurring when $|m| = \sqrt{3}$ in (5.1). In the cases $s = (0; \theta, -\omega)$, $(\lambda_6; \theta, \omega)$, resp. $(\lambda_9; \theta, -1)$, hypothesis (1), (2), resp. (3) of lemma 3.1 is satisfied, by lemma 5.2, 5.3, resp. 5.4. If s is any root with $\langle \rho \,|\, s\rangle = 3$ then it is equivalent to one of these examples under the Γ-stabilizer of ρ, by lemma 5.1. Therefore lemma 3.1 applies to $s^\perp$ for every root s with $\langle \rho \,|\, s\rangle = 3$. It follows that G contains the corresponding loops $\overline{AH}^\circ$. If $\big|\langle \rho \,|\, s\rangle\big| = 3$ then scaling s by a unit reduces to the $\langle \rho \,|\, s\rangle = 3$ case.

The next possible value of $\big|\langle \rho \,|\, s\rangle\big|$ is $2\sqrt{3}$, occurring when $|m| = 2$ in (5.1). In this case lemma 5.2 verifies hypothesis (1) of lemma 3.1, which tells us that G contains $\overline{AH}^\circ$. The next possible value of $\big|\langle \rho \,|\, s\rangle\big|$ is $\sqrt{21}$, occurring when $|m| = \sqrt{7}$. In this case lemma 5.3 verifies hypothesis (2) of lemma 3.1, which tells us that G contains $\overline{AH}^\circ$. The general step of the induction is essentially the same. If $\big|\langle \rho \,|\, s\rangle\big|$ is larger than $\sqrt{21}$, then $|m|$ is larger than $\sqrt{7}$, so lemma 5.2 verifies hypothesis (1) of lemma 3.1. This tells us that G contains $\overline{AH}^\circ$, completing the inductive step. □

References

[A1] D. Allcock, New complex- and quaternionic-hyperbolic reflection groups, *Duke Math. J.* **103** (2000) 303–333.

[A2] D. Allcock, The Leech lattice and complex hyperbolic reflections, *Invent. Math.* **140** (2000) 283–301.

[A3] D. Allcock, The period lattice for Enriques surfaces, *Math. Ann.* **317** (2000) 483–488.

[A4] D. Allcock, A monstrous proposal, in *Groups and Symmetries, From neolithic Scots to John McKay*, ed. J. Harnad. AMS and CRM, (2009). arXiv:math/0606043.

[A5] D. Allcock, On the Y_{555} complex reflection group, *J. Alg.* **322** (2009) 1454–1465.

[ACT1] D. Allcock, J. Carlson and D. Toledo, The complex hyperbolic geometry of the moduli space of cubic surfaces, *J. Algebraic Geom.* **11** (2002) 659–724.

[ACT2] D. Allcock, J. Carlson and D. Toledo, The moduli space of cubic threefolds as a ball quotient, *Mem. Amer. Math. Soc.* **209** (2011). ISBN 978-0-8218-4751-0.
[Ba1] T. Basak, The complex Lorentzian Leech lattice and the bimonster, *J. Alg.* **309** (2007) 32–56.
[Ba2] T. Basak, The complex Lorentzian Leech lattice and the bimonster II, to appear in *Trans. A. M. S.*; arXiv:0811.0062.
[Be] D. Bessis and J. Michel, Explicit presentations for exceptional braid groups, *Experiment. Math.* **13** (2004) 257–266.
[BH] M. Bridson and A. Haefliger, *Metric Spaces of Non-Positive Curvature*, Springer, 1999.
[Br] E. Brieskorn, Die Fundamentalgruppe des Raumes der regulären Orbits einer endlichen komplexen Spiegelungsgruppe, *Invent. Math.* **12** (1971) 57–61.
[C] J. H. Conway, A group of order 8,315,553,613,086,720,000, *Bull. London Math. Soc.* **1** (1969) 79–88.
[CS] J. H. Conway and C. S. Simons, *26 Implies the Bimonster,* J. Alg. **235**, (2001) 805-814.
[CPS] J. H. Conway, R. A. Parker and N. J. A. Sloane, The covering radius of the Leech lattice, *Proc. Roy. Soc. London Ser. A* **380** (1982) 261–290.
[D] I. Dolgachev, Mirror symmetry for lattice polarized K3 surfaces, *J. Math. Sci.* **81** (1996) 2599–2630.
[FN] R. Fox and L. Neuwirth, The braid groups, *Math. Scand.* **10** (1962) 119–126.
[Go] W. Goldman, *Complex hyperbolic geometry* Oxford mathemtical monographs (1999).
[GM] M. Goresky and R. Macpherson *Stratified morse theory*, Springer-Verlag (1980).
[H] G. Heckman, The Allcock Ball Quotient, arXiv:1307.1339.
[HR] G. Heckman and S. Rieken, Two Lorentzian Lattices, preprint.
[HL] G. Heckman and E. Looijenga, The moduli space of rational elliptic surfaces, *Algebraic geometry 2000, Azumino (Hotaka)* pp. 185–248, Adv. Stud. Pure Math. 36, Math. Soc. Japan, Tokyo, 2002.
[Ka] M. Kapovich, *Hyperbolic Manifolds and Discrete Groups*, in Progress in Mathematics 183, Birkhuser, 2001. ISBN 0-8176-3904-7.
[Ko1] S. Kondō, A complex hyperbolic structure for the moduli space of curves of genus three. *J. Reine Angew. Math.* **525** (2000) 219–232.
[Ko2] S. Kondō, The moduli space of curves of genus 4 and Deligne-Mostow's complex reflection groups. *Algebraic geometry 2000, Azumino (Hotaka)*, Adv. Stud. Pure Math., **36**, Math. Soc. Japan, Tokyo, 2002, 383–400.
[L1] E. Looijenga, On the semi-universal deformation of a simple-elliptic hypersurface singularity. Unimodularity. *Topology* **16** (1977), no. 3, 257–262.
[L2] E. Looijenga, The smoothing components of a triangle singularity. II. *Math. Ann.* **269** (1984), no. 3, 357–387.
[L3] E. Looijenga, Artin groups and the fundamental groups of some moduli spaces. *J. Topol.* **1** (2008) 187–216.
[L4] E. Looijenga, The period map for cubic fourfolds. *Invent. Math.* **177** (2009) 213–233.
[LS] E. Looijenga and R. Swierstra, The period map for cubic threefolds. *Compos. Math.* **143** (2007) 1037–1049.
[Na] Y. Namikawa. Periods of Enriques surfaces. *Math. Ann.*, 270:201–222, 1985.
[Ni] V. Nikulin, Finite automorphism groups of Kählerian surfaces of type K3, *Trans. Moscow Math. Soc.* **38** (1980), No 2, 71–135.
[R] J. Ratcliffe, *Foundations of hyperbolic manifolds*. Graduate Texts in Mathematics, 149. Springer-Verlag, 1994. ISBN 0-387-94348-X.
[W] R. Wilson, *The complex Leech Lattice and maximal subgroups of the Suzuki group*. J. Algebra 84 (1983), 151-188.

Department of Mathematics, University of Texas, Austin
E-mail address: allcock@math.utexas.edu
URL: http://www.math.utexas.edu/~allcock

Department of Mathematics, Iowa State University, Ames IA, 50011.
E-mail address: tathastu@gmail.com
URL: https://orion.math.iastate.edu/tathagat/